arXiv:1110.2318v1 [cs.FL] 11 Oct 2011

Compressed Membership for NFA (DFA) with Compressed Labels is in NP (P)

Artur Jeż
University of Wrocław
ul. Joliot-Curie 15, 50-383 Wrocław, Poland
aje@cs.uni.wroc.pl

October 12, 2011

Abstract

In this paper, a compressed membership problem for finite automata, both deterministic and non-deterministic, with compressed transition labels is studied. The compression is represented by straight-line programs (SLPs), i.e. context-free grammars generating exactly one string. A novel technique of dealing with SLPs is introduced: the SLPs are recompressed, so that substrings of the input text are encoded in SLPs labelling the transitions of the NFA (DFA) in the same way, as in the SLP representing the input text. To this end, the SLPs are locally decompressed and then recompressed in a uniform way. Furthermore, such recompression induces only small changes in the automaton, in particular, the size of the automaton remains polynomial.

Using this technique it is shown that the compressed membership for NFA with compressed labels is in NP, thus confirming the conjecture of Plandowski and Rytter [18] and extending the partial result of Lohrey and Mathissen [13]; as it is already known, that this problem is NP-hard, we settle its exact computational complexity. Moreover, the same technique applied to the compressed membership for DFA with compressed labels yields that this problem is in P; for this problem, only trivial upper-bound PSPACE was known.

1 Introduction

1.1 Compression and Straight-Line Programms

Due to ever-increasing amount of data, compression methods are widely applied in order to decrease the data's size. Still, the stored data sometimes need to be processed and decompressing it on each occasion is wasteful, Thus there is a large demand for algorithms working directly on the compressed data. Such task is not as desperate, as it may seem: it is a popular outlook, that compression basically extracts the hidden structure of the text and if the compression rate is high, the text must have a lot of internal structure. And it is natural to assume, that such a structure will help devising methods dealing directly with the compressed representation of the data. Indeed, efficient algorithms for fundamental text operations (pattern matching, checking equality, etc.) are known for various practically used compression methods (LZ, LZW, etc.) [2, 3, 5].

When devising algorithms for compressed data, quite early one needs to focus on the exact compression method, to which the algorithm is applied. The most practical, and challenging, choice is one of the widely used standards, like LZW or LZ. However, a different approach was also proposed: for some applications and for most of theory-oriented considerations it would be useful to *model* one of the practical compression standard by a more mathematically well-founded method. This idea, among other, lay at the foundations of the notion of *Straight-Line Programms* (SLP), whose instance can be simply seen as context-free grammars generating exactly one string.

SLPs are the most popular theoretical model of compression. This is on one hand motivated by a simple, 'clean' and appealing definition, on the other hand, they model the LZ compression standard: each LZ compressed text can be converted into an equivalent SLP with only logarithmic overhead (and in polynomial time) while each SLP can be converted to an equivalent SLP with just a constant overhead (an in polynomial time).

The approach of modelling compression by SLP in order to develop efficient algorithms turned out to be fruitful. Algorithmic problems for SLP-compressed input strings were considered and successfully solved [9, 10, 16]. The recent state-of-the-art efficient algorithms for pattern matching in LZ and LZW compressed text essentially use the reformulation of LZW and LZ methods in terms of SLPs [2, 3]. SLPs found their usage also in programme verification [4, 8]. Attempts were made in order to make out of the box tools for texts represented by SLPs, e.g., efficient indexing structure for SLPs was recently developed [1].

Surprisingly, while SLPs were introduced mainly as a model for practical applications, they turned out to be useful also in strictly theoretical branches of computer science, for instance, their usage was important in the famous proof of Plandowski, that satisfiability of word equations is in PSPACE [17].

1.2 Membership problem

As it should be already clear, that SLPs *are* used, both in theoretical, and applied research in computer science. Hence, tools for them should be developed. In particular, one should be aware, that whenever working with strings, these strings may be supplied as respective SLPs. Hence, all the usual string problems should be reinvestigated in the compressed setting, as the classical algorithms may not apply directly, be inefficient or worse, some of these problems may become computationally difficult.

From language theory point of view, the crucial questions stated in terms of strings, is the one of compressed string recognition. To be more precise, we consider classic membership problems, i.e. recognition by automata, generation by a grammar etc., in which the input is supplied as an SLP. We refer to such problems as *compressed membership problems*. These were first studied in the pioneering work of Plandowski and Rytter [18], who considered compressed membership problem for various formalism for defining languages. Already in this work it was observed, that we should precisely specify, what part of the input is compressed? Clearly the input string, but what about the language representation (i.e. regular expression, automaton, grammar, etc.). Should it be also compressed or not? Both variant of the problem are usually considered, with the following naming convention: when only the input string is compressed, we use a name *compressed membership*, when also the language representation, we prepend *fully* to the name.

In years to come, the compressed membership problem was investigated for various language classes [5, 6, 11, 12, 18]. Compressed word problem for groups and monoids [11, 14, 15], which can be seen as a generalisation of membership problem, was also investigated.

Despite the large attention in the research community, the exact computational complexity of some problems remain open. The most notorious of those is the fully compressed membership problem for NFA, considered already in the work of Plandowski and Rytter [18]. Here, the compression of NFA is done by allowing the transitions by strings, and only by single letters, and representing these strings as SLPs.

It is relatively easy to observe that the compressed membership problem for the NFA is in P, however, the status of the fully compressed variant remained open for a long time. Some partial results were obtained by Plandowski and Rytter [18], who observed that it is in PSPACE and is NP-hard for the case of unary alphabet, both of these bounds being relatively simple. Moreover, they showed that this problem is in NP for some particular cases, for instance, in the case of one-letter alphabet. Further work on the problem was done by Lohrey and Mathissen [13], who demonstrated that if the strings defined by SLP have polynomial periods, the problem is in NP, and when all strings are highly aperiodic, it is in P.

1.3 Our results and techniques

We establish the computational complexity of fully compressed membership problems for both NFAs and DFAs.

Theorem 1. *Fully compressed membership problem for NFA is in* NP*, for DFA it is in* P*.*

Our approach to the problem is essentially different than the approach of Plandowski and Rytter [18] and Lohrey and Mathissen [13]. The main ideas utilised in the previous approaches focused on the properties of strings described by SLPs. We take a completely different approach: we analyse and change the way strings are described by the SLPs in instance. That is, we focus on the SLPs, and not on the encoded strings. Roughly speaking, our algorithm aims at having all the strings in the instance compressed 'in the same way'. To achieve this goal, we decompress the SLPs. Since the compressed text can be exponentially long, we do this locally: we introduce explicit strings into the right-hand sides of the productions. Then, we recompress these explicit strings uniformly. Since such pieces of text are compressed in the same way, we can 'forget' about the original substrings of the input and treat the introduced nonterminals as atomic letters. The idea is that such recompression should shorten the text significantly: roughly one 'round' of recompression, in which every pair of letters that was present at the beginning of the 'round' is compressed, should shorten the encoded strings by a constant factor. It remains to implement the recompression (and changes in the NFA) in npolytime, while keeping the size of N polynomial.

We stress, that the idea of local decompression and synchronous recompression of SLP is new and promising: there is hope that it can be applied to other problem related to SLPs.

2 Preliminaries

2.1 Straight line programmes

Formally, a *Straight line programme* (SLP) is context free grammar G over the alphabet Σ, with a language consisting of exactly one string. Usually it is assumed that G is in a *Chomsky normal form*, i.e. each production is either of the form $X \to YZ$ or $X \to a$. This assumption in particular implies, that strings defined by nonterminals have length at most 2^n; since our algorithm will replace some substrings by shorter ones, none string defined by SLPs during the run of algorithm will exceed this length.

We denote the string defined by nonterminal A by $\mathrm{word}(A)$. This notion extends to $\mathrm{word}(\alpha)$ for $\alpha \in (\mathcal{X} \cup \Sigma)^*$ in the usual way.

2.2 Input

The instance of the fully compressed membership problem for NFA consists of an input string, represented by an SLP, and an NFA N, whose transitions may be labelled by SLPs.

For our purposes it is more convenient to assume, that all SLPs are given as a single context free grammar G with a set of nonterminals $\mathcal{X} = \{X_1, \ldots, X_n\}$, the input string is defined by X_n and the NFAs transitions are labelled with nonterminals of G. Furthermore, in our proof, it is essential to drop the usual assumption, that G is in a Chomsky normal form. However, we still impose some conditions on the productions' right-hand sides. Thus, we require that the grammar during the run of algorithm satisfies the following constraints on its form:

$$\text{each nonterminal has exactly one production, which is of the form} \tag{1a}$$

$$X_i \to uX_jvX_k \quad \text{or} \quad X_i \to uX_jv \quad \text{or} \quad X_i \to u, \quad \text{where } u, v \in \Sigma^* \text{ and } j, k < i, \tag{1b}$$

$$\text{if } \mathrm{word}(X_i) = \epsilon \text{ then } X_i \text{ is not on the right-hand side of any production.} \tag{1c}$$

The strings u, v and their substrings appear *explicitly* in a rule, this notion is introduced to distinguish them from the substrings of $\mathrm{word}(X_i)$. Notice, that (1) does not exclude the case, when $X_i \to \epsilon$ and allowing such a possibility streamlines the analysis.

Without loss of generality we may assume that the input string starts and ends with designated, unique symbols, denoted as \$ and #. These are not essential, however, the first and last letter of word(X_n) need to be treated in a somewhat special manner, furthermore, this applies to their appearances in the NFA as well. Having special symbols for the first and last letter makes the analysis smoother.

2.3 Input size, complexity classes

The size $|G|$ of the representation of grammar G is the sum of length of the right-hand sides of G's rules. The size $|N|$ of the representation of NFA N is the sum of number of its states and transitions. The size $|\Sigma|$ of alphabet Σ is simply the number of elements in Σ.

The input (or, in general, current instance) size is polynomial in N, G, Σ and n, which denotes the number of nonterminals in G. We point out, that one of the crucial properties of our algorithm is that n is not modified during the whole run of the algorithm.

By npolytime (polytime) we denote the class of algorithms running in non-deterministic (deterministic, respectively) polynomial time, and by NP (P, respectively) the corresponding complexity classes of the decision problems.

2.4 Known results

We use the following basic result, which states that the fully compressed membership problem, when the input string is over a unary alphabet, is in NP for NFA and in P for DFA.

Lemma 1 (cf. [18, Theorem 5])**.** *The fully compressed membership problem restricted to the input string over an alphabet $\Sigma = \{a\}$ is in* NP *for NFA and in* P *for DFA.*

The first claim can be easily inferred from the result of Plandowski and Rytter [18, Theorem 5], who proved that this problem is in NP, when $\Sigma = \{a\}$, i.e. also transitions in the NFA are labelled by powers of a only. The second claim is trivial.

2.5 Path and labels in NFA

Since we deal with automata, proofs will be concerned with (accepting) paths for strings. We shall consider NFAs, for which transitions are labelled with either letters, or non-terminals of G. That is, that transition relation δ satisfies $\delta \subseteq Q \times (\Sigma \cup \mathcal{X}) \times Q$. Consequently, a *path* $\mathcal{P}$ from state p_1 to p_{k+1} is a sequence $\alpha_1\alpha_2 \ldots \alpha_k$, where $\alpha_i \in \Sigma \cup \mathcal{X}$ and $\delta(q_i, \alpha_i, q_{i+1})$. We write, that $\mathcal{P}$ *induces* such a list of labels. The word($\mathcal{P}$) *defined* by such a path $\mathcal{P}$ is simply word($\alpha_1 \ldots \alpha_k$). We also say that $\mathcal{P}$ is a path for a string word($\mathcal{P}$). A path is *accepting*, if it ends in an accepting state. A string w is accepted by N if there is an accepting path from the starting state for w.

3 Basic classifications and outline of the algorithm

In this section we present the outline of the algorithm for fully compressed membership for NFAs. Its main part consist of recompression, i.e. replacing strings appearing in word(X_n) by shorter ones. In some cases, such replacing is harder, in other easier. It should be intuitively clear, that it depends on the position of letters inside encoded texts: if a is a first or last letter of some word(X_i), then recompressing strings including a looks difficult, as such strings can be split into different nonterminals and recompression requires heavy modification of G, or even rebuilding of the NFA. On the other hand, if a letter a is only 'inside' strings encoded by nonterminals, its compression is done only 'inside' rules of G, which seems easy. Thus, before we state the algorithm, we firstly introduce classification of letters (and strings) into 'easy' and 'difficult' to compress.

3.1 Crossing appearances, types of letters, non-extendible appearances

We say that a string w has a *crossing appearance in a (string defined by) nonterminal* X_i with a production $X_i \to uX_jvX_k$, if w appears in word(X_i), but this appearance is not contained in neither u, v, word(X_j) nor word(X_k). Intuitively, this appearance 'crosses' the symbols in of $u\,\text{word}(X_j)v\,\text{word}(X_k)$,i.e, at the same time part of w is in the explicit substring (u or v) and part is in the compressed strings (word(X_j) or word(X_k)). This notion is similarly defined for nonterminals with productions of the form $X_i \to uX_jv$, productions of the form $X_i \to u$ clearly do not have crossing appearances.

A string w has a *crossing appearance in the NFA* N, if there is a path in N inducing list of labels $\alpha_1\alpha_2$, where $\alpha_1, \alpha_2 \in \{X_1, \ldots, X_n\} \cup \Sigma$ with at least one of α, α_2 being a nonterminal, such that w appears in a word$(\alpha\alpha_2)$, but this appearance is not contained in the word(α_1), nor in word(α_2). The intuition is similar as in the case of crossing appearance in a rule: it is possible that a string w is split between two transitions' labels. Still, there is nothing difficult in consecutive letter transitions, thus we treat such a case as a simple one.

We say that a pair of different letters ab is a *crossing pair*, if ab has a crossing appearance of any kind. Otherwise, such a pair is *non-crossing*.

We say that a letter $a \in \Sigma$ is *right-outer* (*left-outer*), if there is a nonterminal X_i, such that a is the leftmost (rightmost, respectively) symbol in word(X_i). A letter is *outer*, if it is left-outer or right-outer. Otherwise, the letter is *inner*. Notice, that if a pair ab is crossing, then a is left-outer or b is right-outer. The outer letters and crossing pairs correspond to the intuitive notion of 'hard' to compress.

The following lemma shows, that while G may encode long strings, they have relatively few different short substrings and few outer letters.

Lemma 2. *There are at most* $2n$ *different outer letters and at most* $|G| + 3n$ *different pairs of letters appearing in* word(X_n), ..., word(X_n).

The set of outer letters, the set of crossing pairs and the set of non-crossing pairs appearing in word(X_n), ..., word(X_n) *can be computed in* polytime.

The notions of (non-) crossing pairs do not apply to aa, still, an analog can be defined: for a letter $a \in \Sigma$ we say that a^ℓ is a a's *non-extendible appearance* of length ℓ, if it appears in some string defined by some nonterminal and it is surrounded by letters other than a, formally, if there exist two letters $x, y \in \Sigma$, where $x \neq a \neq y$ and a nonterminal X_i, such that $xa^\ell y$ is a substring of word(X_i). Similarly to crossing pairs, it can be shown that there are not too many different non-extendible appearances of a.

Lemma 3. *For an inner letter* a *and a grammar* G *there are at most* $|G|$ *different lengths of* a*'s non-extendible appearances in* word(X_1), ..., word(X_n). *The set of these lengths can be calculated in* polytime.

3.2 Outline of the algorithm

Our algorithm consists of two main operations performed on strings encoded by G

appearance compression of a For each a^ℓ that has a non-extendible appearance in word(X_n), replace all a^ℓs in word(X_1), ..., word(X_n) by a fresh letter a_ℓ. Modify N accordingly.

pair compression of ab For two *different* letters ab replace each of ab in word(X_1), ..., word(X_n) by a fresh letter c. Modify N accordingly.

We denote the string obtained from w by compression of appearances of a by $AC_a(w)$, and the string obtained by compression of a pair ab into c by $PC_{ab\to c}(w)$.

We adopt the following notational convention throughout rest of the paper: whenever we refer to a letter a_ℓ, it means that the last appearance compression was done for a and a_ℓ is the letter that replaced a^ℓ.

The main idea behind the algorithm is that appearance compression and pair compression shorten the encoded texts significantly. The challenging part of the algorithm is the modification of NFA N 'accordingly' to the changes of SLPs. The general schema is given by Algorithm 1.

Algorithm 1 Outline of the main algorithm

```
1: while |word(X_n) > n| do
2:     while something changed do
3:         for a: inner letter do
4:             compress appearances of a, modify N accordingly
5:         for non-crossing pair ab in word(X_n), a, ∉ {$, #} do
6:             compress ab, modify N accordingly
7:     L ← list of outer letters, except $ and #
8:     for a ∈ L do
9:         compress appearances of a, modify N accordingly
10:        for each a_ℓb in word(X_n) do
11:            compress a_ℓb, modify N accordingly
12: Decompress X_n and solve the problem naively.
```

There are two important remarks to be made:

- there is no explicit non-deterministic operation in the code, however, it appears implicitly in the term 'modify the NFA accordingly' in lines 4 and 9. Roughly, one needs to solve fully compressed membership for string a^ℓ for that, and this is known to be NP-hard.
- the compression (both of pairs and appearances) is never applied to \$, nor to #. The markers were introduced so that we do not bother with strange behaviour where first or last letter is compressed, and so we do not touch the markers.

It should be more or less obvious, that the compressions performed by Algorithm 1 shortens word(X_n). This is formally stated as follows.

Lemma 4. *There are $\mathcal{O}(n)$ executions of the loop in line 1 of Algorithm 1.*

Remark 1. Notice, that pair compression $PC_{ab\to b}$ is in fact introducing a new nonterminal with a production $c \to ab$, similarly AC_a. Hence, Algorithm 1 creates new SLPs, encoding strings from the instance. However, these new nonterminals are never expanded, they are always treated as individual symbols. Thus it is better to think of them as letters. In particular, the analysis of running time of Algorithm 1 relies on the fact, that no new nonterminals are introduced by Algorithm 1.

4 Details

In this section we describe in detail how to implement the appearance compression and pair compression. In particular, we are going to formulate the connections between NFA and SLPs preserved during Algorithm 1 and demonstrate how to modify the NFA accordingly.

4.1 Invariants

The invariants below describe the connection between the grammar kept by Algorithm 1 and the input one.

SLP 1 The set of used nonterminals is a subset of $\mathcal{X} = \{X_1, \ldots, X_n\}$ and the productions are of the form described in (1).

SLP 2 For every production $X_i \to \alpha_i$, the original instance contained a production $X_i \to \alpha'_i$, where the nonterminals appearing in α appear (in the same order) in α'.

SLP 3 The nonterminal X_n has a production $X_n \to \$uX_{n-1}v\#$, where $u, v \in (\Sigma \setminus \{\$, \#\})^*$; \$, # are not used in other productions.

The following invariants represent the constraints on the NFA.

Aut 1 every transition of N is labelled by a single letter of Σ (*letter transition*) or by a nonterminal (*nonterminal transition*), each nonterminal labels at most one transition. No transition is labelled with X_n.

Aut 2 there is a unique starting state that has a unique outgoing transition, by letter \$, and no incoming transitions; there is no other transition by \$. Similarly, there is a unique accepting state that has a unique incoming transition, by letter #, it does not have any outgoing transitions; there is no other transition by # in N.

Algorithm 1 will preserve (SLP 1)–(Aut 2), and we shall always assume, that the input of the subroutines satisfies (SLP 1)–(Aut 2).

We assume that the input instance satisfies (SLP 1)–(Aut 2), moreover, that that the input grammar is in the Chomsky normal form. It is routine to transform (in polytime) the input instances not satisfying these conditions into equivalents instance that satisfy them.

4.2 Compression of non-crossing pairs and inner letters

The compression of non-crossing pairs and appearance compression for inner letters is intuitively easy: whenever these appear in strings encoded by G or on paths in N, they cannot be split between nonterminals or between transitions. Thus, it should be enough to replace their explicit appearances in the grammar and in the NFA. This is formalised and shown in this subsection

4.2.1 Compression of non-crossing pairs

We first demonstrate, how to perform the pair compression for non-crossing pairs. Consider a non-crossing pair ab. Since it is non-crossing, it can only appear in the the explicit strings in the rules of G. Hence, compressing ab into a fresh letter c consists simply of replacing each explicit ab by c in right-hand side of every production. Still, ab can appear on a path in N. But since ab is non-crossing, this can be either wholly inside a nonterminal transition (and so compression was already taken care of), or on two consecutive letter transitions. This is also easy to handle: whenever there is a path from p to q by a string ab, we introduce a new letter transition by c from p to q. This description is formalised in Algorithm 2.

To distinguish between the input and output G and N, we utilise the following convention: 'unprimed' names refer to the input (like G, X_i, N), while 'primed' symbols refer to the output (like G, X'_i, N'). This convention is used for lemmata concerning algorithms through the paper.

Algorithm 2 Pair compression for a non-crossing pair ab

```
1: for each production X_i → α do                    ▷ Pair compression
2:     replace each explicit ab in α by c
3: for states p, q do                                ▷ Appropriate modifications of N
4:     if there is a path for ab from p to q then
5:         put a transition δ_N(p, c, q)
```

We show, that Algorithm 2 indeed realises the pair compression for a non-crossing pair.

Lemma 5. *Algorithm 2 runs in* polytime *and preserves (SLP 1)–(Aut 2). When applied to a non-crossing pair of letters ab, where* $a, b \notin \{\$, \#\}$*, it implements the pair compression, i.e.* $\mathrm{word}(X'_i) = PC_{ab\to c}(\mathrm{word}(X_i))$*, for each* X_i*.*

N' *recognises* $\mathrm{word}(X'_n)$ *if and only if* N *recognises* $\mathrm{word}(X_n)$*. If* N *is a DFA, so is* N'*.*

4.2.2 Inner letter appearance compression

We can apply the same approach, as the one used for non-crossing pairs compression, to the inner letters appearance compression. However, in this case, the modification of N uses non-determinism.

Since a is an inner letter, it cannot appear as the last or first letter of any nonterminal, and so every non-extendible appearance of a in $\mathrm{word}(X_1), \ldots, \mathrm{word}(X_n)$ is an explicit substring in the right-hand sides of G's rules; so we simply replace explicit a^ℓ by a fresh letter a_ℓ in each right-hand side of G's rules. Before considering the NFA, notice that as a is an inner letter, a^ℓ cannot have a crossing appearance in N, and no nonterminal defines a^ℓ. Hence, when a^ℓ is a substring of a string defined by a path in N, then a^ℓ appears wholly inside a nonterminal transition, and this was already taken care of when considering G, or a^ℓ labels a path using letter transitions only. So it is enough to check, whether there is a path for a^ℓ from p to q using only letter transitions.

Algorithm 3 Appearance compression for an inner letter a

1: establish the lengths $\ell_1, \ldots, \ell_k$ of a's non-extendible appearance
2: **for** each a^{ℓ_m} **do** ▷ Appearance compression
3: **for** each production $X_i \to \alpha$ **do**
4: replace every explicit non-extendible a^{ℓ_m} in in α by a_{ℓ_m}
5: **for** states p, q in N **do** ▷ Appropriate modifications of N
6: **if** $\delta_N(p, a^{\ell_m}, q)$ **then** ▷ Verify non-deterministically, see Lemma 1
7: put a transition $\delta_N(p, a_{\ell_m}, q)$

Lemma 6. *Suppose that Algorithm 3 is applied for inner letter $a \notin \{\$, \#\}$. Then Algorithm 3 properly implements non-extendible appearance compression, i.e.* $\mathrm{word}(X_i') = AC_a(\mathrm{word}(X_i))$ *for each X_i and preserves (SLP 1)–(Aut 2).*

The operations in line 6 of Algorithm 3 can be performed in npolytime*, other operations can be performed in* polytime*.*

Each of the new letters a_ℓ is inner. N recognises $\mathrm{word}(X_n)$ *if and only if N' recognises* $\mathrm{word}(X_n')$ *for some non-deterministic choices. If N is DFA, so is N'.*

4.3 Compression of outer letters and crossing pairs

Now, we turn our attention to the compression of outer letters and crossing pairs. The outline is as follows: we fix an outer letter a and modify the instance, so that a becomes inner. Then, Algorithm 3 is applied to a. Next, we want to compress each pair of the form $a_\ell b$. Such a pair can be crossing, as b can be a left-outer letter. Thus, we modify the instance again, so that none of $a_\ell b$ is a crossing pair so that it can be compressed using Algorithm 2.

4.3.1 Transforming an inner letter to an outer letter

The reason, why a is an outer letter, is that it is the first or the last symbol in some $\mathrm{word}(X_i)$. To make it an inner letter, it is enough to remove each nonterminal's a-prefix and a-suffix. To be more precise: fix i and let $\mathrm{word}(X_i) = a^{\ell_i} u a^{r_i}$, where u does not start nor end with a. Then our goal is to modify G so that $\mathrm{word}(X_i') = u$. (If $\mathrm{word}(X_i)$ is a power of a, we simply give $u' = \epsilon$ and $r_i = 0$.) This can be done in a bottom-up fashion, starting from X_1: it is enough to calculate and memorise the lengths of the a-prefixes and a-suffixes for consecutive nonterminals, see the loop in lines 1–4 of Algorithm 4. Then we need to modify the NFA accordingly: it is enough to replace the transition labelled with X_i by path consisting of three transitions, labelled with a^{ℓ_i}, X_i' and a^{r_i}.

Algorithm 4 Changing an inner letter a to an outer letter.

1: **for** $i = 1 \,..\, n$ **do** ▷ Removing a-prefix and suffix
2: let the production for X_i be $X_i \to \alpha_i$
3: replace any nonterminal X_j in α_i by $a^{\ell_j} X_j a^{r_j}$, remove those defining ϵ
4: calculate the explicit a-prefix a^{ℓ_i} and a-suffix a^{r_i} of α_i and remove them
5: **if** there is a nonterminal transition $\delta_N(p, X_i, q)$ in N **then** ▷ modification of NFA
6: create new states p_1, q_1 in N, remove transition $\delta_N(p, X_i, q)$
7: set transitions: $\delta_N(p, a^{\ell_i}, p_1)$, $\delta_N(p_1, X_i, q_1)$, $\delta_N(q_1, a^{r_i}, q)$

The removed a-prefixes and a-suffixes can be exponentially long, and so we store them in the rules in a succinct way, i.e. a^ℓ is represented as (a, ℓ); the size of representation of ℓ is $\mathcal{O}(\log \ell)$, that is, linear in n. We say, that such a grammar is in an *a-succinct* form. The situation is similar for the NFA, as it might have transitions labelled with a^ℓ, which are stored in succinct way as well. We say that N satisfies *a-relaxed* (Aut 1), if its transitions are labelled by nonterminals, a single letter or by a^ℓ, where $\ell \leq 2^n$.

Lemma 7. *Algorithm 4 applied to letter $a \notin \{\$, \#\}$ runs in* polytime *time and preserves (SLP 1)–(Aut 2), except that it a-relaxes (Aut 1). G' is in the a-succinct form.*

Let $\mathrm{word}(X_i) = a^{\ell_i} u_i a^{r_i}$*, where u_i does not begin, nor end with a. Then (after running Algorithm 4),* $\mathrm{word}(X'_i) = u_i$*. In particular, after running Algorithm 4 the letter a is inner.*

N accepts $\mathrm{word}(X_n)$ *if and only if N' accepts* $\mathrm{word}(X'_n)$*. If N is a DFA, so is N'.*

Since a is no longer an outer letter, we may compress its non-extendible appearances using Algorithm 3. Some small twitches are needed to accommodate the a-succinct form of G and the fact that N is a-relaxed, though basically we just perform appearances compression for inner a, as described by Algorithm 3: the non-trivial part of Algorithm 3 was the application of Lemma 1, which works for such large powers of a in npolytime. Other actions of Algorithm 3 generalise in a simple way.

Lemma 8. *Algorithm 3 can be extended, so that it applies to instances satisfying (SLP 1)–(SLP 3) with G in the a-succinct form and a-relaxed-(Aut 1)–(Aut 2). Lemma 6 applies to such an extension. The output satisfies (SLP 1)–(Aut 2).*

4.3.2 Crossing pair compression

By Lemma 8 all letters a_ℓ are inner. However, a pair of the form $a_\ell b$ can still be crossing and this can happen only when b is a left-outer letter, so we would like to make such b not a left-outer letter. To do so, we 'pop' one letter from the beginning of each nonterminal (that is, all left outer letters), i.e. we modify the grammar so that $\mathrm{word}(X_i) = \mathrm{first}(X_i)\,\mathrm{word}(X'_i)$, where $\mathrm{first}(X_i)$ denotes the first letter of $\mathrm{word}(X_i)$. Clearly, after such operations there are some (perhaps other) left-outer letters in G. Still, we show, that none $a_\ell b$ is crossing.

Popping letters is performed similarly to the removal of the a-prefix, i.e. in a bottom-up fashion, starting from X_1: in a rule $X_i \to \alpha$ it is enough to replace each nonterminal X_j in α by $\mathrm{first}(X_j)X'_j$, then store the first letter from α in $\mathrm{first}(X_i)$ and remove it from α. It is easy to modify the NFA N accordingly: when there is a transition $\delta_N(p, X_i, q)$, we change it into a chain of two transitions: $\delta_{N'}(p, \mathrm{first}(X_i), p_1)$ and $\delta_{N'}(p_1, X'_i, q)$. This operation is not performed on X_n, as the letter \$ is not going to be compressed anyway. There is a little detail to take care of: if $|\,\mathrm{word}(X_i)| = 1$ then popping a letter from X_i creates X'_i, which defines ϵ. Then X'_i should be removed from the right-hand sides of the rules and in the NFA N we simply replace the transition by X_i by a transition by $\mathrm{word}(X_i)$. This description is formalised in Algorithm 5.

Algorithm 5 Popping letters

for $i \leftarrow 1 \,..\, n-1$ **do** ▷ Popping letters
  let $X_i \to \alpha$
  replace each X_j in α by $\text{first}(X_j)X_j$, remove those defining ϵ
  set $\text{first}(X_i) \leftarrow$ the first letter of $\text{word}(X_i)$
  if $i < n$ **then**
    remove the first letter from α
  if there is a transition $\delta_N(p, X_i, q)$ in N **then** ▷ NFA modification
    remove transition $\delta_N(p, X_i, q)$
    if $|\,\text{word}(X_i)| > 1$ **then**
      create new state p_1 in N, set transitions: $\delta_N(p, \text{first}(X_i), p_1)$, $\delta_N(p_1, X_i, q)$
    else
      set transition $\delta_N(p, \text{first}(X_i), q)$

Lemma 9. *Algorithm 5 runs in time* polytime *and preserves (SLP 1)–(Aut 2). Let* $\text{word}(X_i) = bu$*, where* $b \in \Sigma$*, then* $\text{word}(X_i') = u$ *for* $i < n$ *and* $\text{word}(X_n') = \text{word}(X_n)$*. After running Algorithm 5, pairs of the form* $a_\ell b$ *appearing in* $\text{word}(X_n)$ *are non-crossing.*

N' *accepts* $\text{word}(X_n')$ *if and only if* N *accepts* $\text{word}(X_n)$*. If* N *is deterministic, so is* N'*.*

Now, it is enough to apply the pair compression for non-crossing pairs to each pair of the form $a_\ell b$. For convenience, we write the whole procedure for pair compression for crossing pairs in Algorithm 6.

Algorithm 6 Pair compression for crossing pairs $a_\ell b$

pop the first letter from each nonterminal (run Algorithm 5)
for each a_ℓ **do**
  for each b such that $a_\ell b$ appears in $\text{word}(X_n)$ **do**
    run Algorithm 2 for a pair $a_\ell b$

Lemma 10. *Algorithm 6 runs in* polytime *and preserves (SLP 1)–(Aut 2). It implements pair compression for* ab*, in the sense that* $\text{word}(X_i') = PC_{ab\to c}(\text{word}(X_i))$ *for each* X_i*.*

N' *accepts* $\text{word}(X_n')$ *if and only if* N *accepts* $\text{word}(X_n)$*. If* N *is deterministic, so is* N'*.*

4.4 Running time

Since the running time of each algorithm is npolytime, it is enough to show that the size of Σ, G and N are always polynomial in n (recall, that n is unchanged throughout Algorithm 1).

Lemma 11. *During Algorithm 1, the sizes of* Σ*,* G*,* N *are polynomial in* n*.*

Using Lemmas 5–11 it is now possible to conclude that Algorithm 1 correctly solves the fully compressed membership problem for NFA, in nondeterministic polynomial (in n) time. The only source of non-determinism is the one in Lemma 1, and so for DFA the corresponding problem can be solved deterministically.

Acknowledgements

I would like to thank Paweł Gawrychowski for introducing me to the topic and for pointing out the relevant literature.

References

[1] Philip Bille, Gad M. Landau, Rajeev Raman, Kunihiko Sadakane, Srinivasa Rao Satti, and Oren Weimann. Random access to grammar-compressed strings. In Randall [19], pages 373–389.

[2] Paweł Gawrychowski. Optimal pattern matching in LZW compressed strings. In Randall [19], pages 362–372.

[3] Paweł Gawrychowski. Pattern matching in Lempel-Ziv compressed strings: fast, simple, and deterministic. In Camil Demetrescu and Magnús M. Halldórsson, editors, *ESA*, volume 6942 of *LNCS*, pages 421–432. Springer, 2011.

[4] Blaise Genest and Anca Muscholl. Pattern matching and membership for hierarchical message sequence charts. *Theory of Computing Systems*, 42(4):536–567, 2008.

[5] Leszek Gąsieniec, Marek Karpiński, Wojciech Plandowski, and Wojciech Rytter. Efficient algorithms for Lempel-Ziv encoding. In Rolf G. Karlsson and Andrzej Lingas, editors, *SWAT*, volume 1097 of *LNCS*, pages 392–403. Springer, 1996.

[6] Artur Jeż and Alexander Okhotin. One-nonterminal conjunctive grammars over a unary alphabet. *Theory of Computing Systems*, 49(2):319–342, 2011.

[7] Rastislav Královič and Paweł Urzyczyn, editors. *Mathematical Foundations of Computer Science 2006, 31st International Symposium, MFCS 2006, Stará Lesná, Slovakia, August 28–September 1, 2006, Proceedings*, volume 4162 of *LNCS*. Springer, 2006.

[8] Sławomir Lasota and Wojciech Rytter. Faster algorithm for bisimulation equivalence of normed context-free processes. In Královič and Urzyczyn [7], pages 646–657.

[9] Yury Lifshits. Solving classical string problems an compressed texts. In Rudolf Ahlswede, Alberto Apostolico, and Vladimir I. Levenshtein, editors, *Combinatorial and Algorithmic Foundations of Pattern and Association Discovery*, volume 06201 of *Dagstuhl Seminar Proceedings*. IBFI, Schloss Dagstuhl, Germany, 2006.

[10] Yury Lifshits and Markus Lohrey. Querying and embedding compressed texts. In Královič and Urzyczyn [7], pages 681–692.

[11] Markus Lohrey. Word problems and membership problems on compressed words. *SIAM Journal of Computing*, 35(5):1210–1240, 2006.

[12] Markus Lohrey. Compressed membership problems for regular expressions and hierarchical automata. *International Journal of Foundations of Computer Sciences*, 21(5):817–841, 2010.

[13] Markus Lohrey and Christian Mathissen. Compressed membership in automata with compressed labels. In Alexander S. Kulikov and Nikolay K. Vereshchagin, editors, *CSR*, volume 6651 of *LNCS*, pages 275–288. Springer, 2011.

[14] Markus Lohrey and Saul Schleimer. Efficient computation in groups via compression. In Volker Diekert, Mikhail V. Volkov, and Andrei Voronkov, editors, *CSR*, volume 4649 of *LNCS*, pages 249–258. Springer, 2007.

[15] Jeremy MacDonald. Compressed words and automorphisms in fully residually free groups. *International Journal of Automation and Computing*, 20(3):343–355, 2010.

[16] Wojciech Plandowski. Testing equivalence of morphisms on context-free languages. In Jan van Leeuwen, editor, *ESA*, volume 855 of *LNCS*, pages 460–470. Springer, 1994.

[17] Wojciech Plandowski. Satisfiability of word equations with constants is in PSPACE. *Journal of the ACM*, 51(3):483–496, 2004.

[18] Wojciech Plandowski and Wojciech Rytter. Complexity of language recognition problems for compressed words. In Juhani Karhumäki, Hermann A. Maurer, Gheorghe Paun, and Grzegorz Rozenberg, editors, *Jewels are Forever*, pages 262–272. Springer, 1999.

[19] Dana Randall, editor. *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23–25, 2011.* SIAM, 2011.

Appendix

A Additional material for Section 1

(see page 3)

proof of Theorem 1. The proof follows by showing that Algorithm 1 properly verifies, whether word(X_n) is accepted by N and that Algorithm 1 runs in npolytime.

Let us first show correctness of Algorithm 1. All subroutines of Algorithm 1 (non-deterministically) modify the instance, changing G, N and X_n into G', N' and X'_n (notice, that the output depends on the non-deterministic choices). Let $N^{(i)}$, $G^{(i)}$, $X_n^{(i)}$ for $i = 1, \ldots, k$ be the consecutive obtained instances, with $i = 1$ representing the input instance. Then $N^{(i)}$ accepts word$(X_n)^{(i)}$ if and only if for some non-deterministic choices the resulting $N^{(i+1)}$ accepts word$(X_n)^{(i+1)}$. This is shown in Lemmata 5, 6, 7, 8, 9, 10 (if some of the procedures are deterministic, then the output does not depend on any choices). So, if $N^{(1)}$ does not accept word$(X_1^{(1)})$ also $N^{(k)}$ does not accept word$(X_1^{(k)})$. On the other hand, if $N^{(1)}$ accepts word$(X_n^{(1)})$, then there exits a sequence of instances (representing proper non-deterministic guesses), such that for each i $N^{(i)}$ accepts word$(X_n^{(i)})$. In particular, $N^{(k)}$ does accept word$(X_n^{(k)})$ and as $|\,\text{word}(X_1^{(k)})| < n$, word$(X_1^{(k)})$ can be decompressed and acceptance by $N^{(k)}$ can be checked naively in polytime.

Now, we should show that the running time is in fact (non-deterministic) polynomial. Lemmata 5, 6, 7, 8, 9, 10 claim, that each of the subroutine runs in time npolytime in the size of the current instance. However, by Lemma 11, the size of this instance is polynomial in n. So, it is left to show that each of the subroutine is run only polynomially many times. Notice, that each invocation of Algorithms 2, 3, 6 introduces a new letter to Σ, and we know by Lemma 11 that the final size of Σ is poly(n). Moreover, each invocation of Algorithm 4 if followed by Algorithm 3 (in the extended version), similarly, each Algorithm 5 can be associated with Algorithm 6 invoked after it. And so also Algorithms 4, 5 are invoked poly(n) times. So the whole running time is npolytime.

It is left to show, that if the input is a DFA, Algorithm 1 can be determinised. Firstly, notice that by Lemmata 5, 6, 7, 8, 9, 10, if the instance consisted of a DFA, each instance kept by Algorithm 1 is also a DFA. Observe, that the only non-deterministic choices in Algorithm 1 are performed when calling a subroutine for a fully compressed membership problem for a string over an alphabet consisting of a single letter (see Lemma 1). However, the same lemma states, that when the input consists of a deterministic automaton, the problem is in P. Thus, there is no non-determinism in Algorithm 1. □

B Additional material for Section 2

It is more convenient to represent path's list of labels as

$$u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n},$$

where each $u_{i_j} \in \Sigma^*$ is a string representing the consecutive letter labels and X_{i_j} represents a nonterminal label. Notice, that u_{i_j} may be empty. We write, that $\mathcal{P}$ *induces* such a list of labels. The word$(\mathcal{P})$ *defined* by such a path $\mathcal{P}$ is

$$\text{word}(\mathcal{P}) = u_{i_1}\,\text{word}(X_{i_2})u_{i_3}\,\text{word}(X_{i_4})\cdots\text{word}(X_{i_{n-1}})u_{i_n}.$$

(see page 4)

Proof of Lemma 1. Notice, that if the input string w is over an alphabet $\{a\}$, no accepting path in NFA for w may use transitions that denote strings having letters other than a. Thus, any such transitions can be deleted from N and we end up with an instance, to which the result of Plandowski and Rytter [18, Theorem 5] can be applied directly.

Consider now the deterministic automaton. As shown in the previous paragraph, we can limit ourselves to transitions by powers of a. Since the automaton is deterministic, for each state there

is at most one transition labelled with a power of a, and so the path for w cycles after at most n transitions. As the length of the cycle can be calculated in polytime, the whole problem can be easily checked in polytime. □

C Additional material for Section 3

(see page 5)

Proof of Lemma 2. Since there are n nonterminals, there are at most n left-outer letters and n right-outer letters. Clearly, they can be calculated in polytime.

We estimate the total number of pairs of letters that appear in $\mathrm{word}(X_1), \ldots, \mathrm{word}(X_n)$. Consider first such pairs that appear in some explicit string on the right-hand side of some production. Since the total length of explicit strings is $|G|$, there are at most $|G|$ such pairs. Other pairs are assigned to nonterminals $X_1, \ldots, X_n$: a pair ab is assigned to X_i, if it appears in $\mathrm{word}(X_i)$ and it does not appear in $\mathrm{word}(X_1), \ldots, \mathrm{word}(X_{i-1})$. We show, that at most three different pairs are assigned to each nonterminal. In this way, total number of different crossing-pairs is at most $|G|+3n$. Indeed, if ab is assigned to X_i with a production $X_i \to uX_jvX_k$, then ab does not appear neither in u, v, as in such case it was already accounted; nor in $\mathrm{word}(X_j)$, $\mathrm{word}(X_k)$, as in such case it is not assigned to X_i. Thus, there are three possibilities:

- a is the last letter of u and b is the first letter of $\mathrm{word}(X_j)$,
- a is the last letter of $\mathrm{word}(X_j)$ and b is the first letter of v,
- a is the last letter of v and b is the first letter of $\mathrm{word}(X_k)$.

The cases, in which u or v is empty or there are less nonterminals on the right-hand side of the production, are similar.

The above description can be turned to a straightforward algorithm computing both the list of all non-crossing and crossing pairs appearing in $\mathrm{word}(X_1), \ldots, \mathrm{word}(X_n)$. First, the list of all pairs of letters with such appearances is calculated: clearly, it is enough to read every rule (for X_i) and store the pairs that appear in the explicit strings and the pairs that are assigned to X_i. Then, for each pair of letters it should be decided, whether it is crossing. To this end, we check, whether it has a crossing appearance in any nonterminal or in N, which can be done in polytime. Such pairs are crossing, other are non-crossing. □

Instead of Lemma 3 we show an extended version

Lemma 12 (stronger variant of Lemma 3)**.** *For an inner letter a and a grammar G, which can be given in an a-succinct form, there are at most $|G|$ different lengths of a's non-extendible appearances in* $\mathrm{word}(X_1), \ldots \mathrm{word}(X_n)$*. The set of these lengths can be calculated in* polytime*.*

(see page 5)

Proof. The proof is similar to the proof of Lemma 2.

Since a is an inner letter, all of its non-extendible appearances are explicit substrings in productions of G. In particular, each symbol in the productions can uniquely assigned to the non-extendible appearance to which it belongs; this is true regardless of whether the symbol represents letter or block of letters written in a succinct way. Thus, there are at most $|G|$ different non-extendible appearances for a. To calculate the lengths of these appearances, it is enough to read the explicit strings in the rules, adding the appropriate lengths, which can be done in polytime, as these lengths are at most 2^n. □

(see page 6)

Proof of Lemma 4. Consider any 2 consecutive letters ab, where $a \neq \$$ and $b \neq \#$, appearing in the $\mathrm{word}(X_n)$ at the beginning of loop starting in line 1. We show, that at least one of these two letters is compressed before the next execution of this loop. In this way, if we partition $\mathrm{word}(X_n)$ into blocks of 4 consecutive letters, each block is shortened by at least one letter in each iteration of the loop from line 1. Thus the length of $\mathrm{word}(X_n)$ decreases by a factor of $3/4$ in each iteration and so this loop is executed at most $\mathcal{O}(n)$ times, as in the input instance $|\,\mathrm{word}(X_n)| \leq 2^n$.

Assume for the sake of contradiction, that none of letters a, b is compressed during this iteration of the loop.

If $a = b$, then this pair of consecutive letters is going to be compressed, either in line 4 or 9, depending on whether a is inner or outer. Contradiction.

So suppose now that $a \neq b$. Since the pair ab was not compressed before line 7 in Algorithm 1, it is crossing in this line and thus at least one of the letters a, b is outer. This outer letter is going to be compressed with the next letter in line 11 (or earlier). Contradiction.

This ends the case inspection.

Since in the input instance it holds that $P \leq 2^n$, this guarantees, that the loop in line 1 of Algorithm 1 is run $\mathcal{O}(n)$ times. □

D Additional material for Section 4

(see page 7)

Proof of Lemma 5. The bound on the running time is obvious from the code.

Since Algorithm 2 only modifies the grammar by shortening some strings in the productions (it does not create ϵ-rules), and it does not affect \$ and # in the rules, the (SLP 1)–(SLP 3) are preserved. The only modification in N is the introduction of new transition by a single letter (namely, by c) between states that are joined by a path for ab. The only change in N is the introduction of a new letter transitions. Moreover, if there is a new transition $\delta_{N'}(p', c, q')$, then p has an outgoing production by $a \notin \{\$, \#\}$, and so it was not a starting or accepting state, and q had an incoming transition by $b \notin \{\$, \#\}$, and similarly it was not a starting nor accepting state. Thus (Aut 1)–(Aut 2) hold for N' as well. Notice, that if N is deterministic, so is N': suppose that there are two different transitions by a letter d from state p in N'. If $d \neq c$, then these two transition are also present in N, which is not possible, as N is deterministic. If $d = c$, then in N there are two different paths from p for a string ab, which is also a contradiction.

We now show that N' recognises word(X'_n) if and only if N recognises word(X_n). To this end we demonstrate, how Algorithm 2 affects word(X_i):

Claim 1. After performing Algorithm 2, it holds that

$$\text{word}(X'_i) = PC_{ab\to c}(\text{word}(X_i)). \tag{2}$$

Proof. Notice, that as $a \neq b$, $PC_{ab\to c}$ is well defined for each string.

The claim follows by a simple induction on the nonterminal's number: Indeed, this is true when the production for X_i has no nonterminal on the right-hand side (recall the assumption that $a \neq b$), as in this case the pair compression on right hand side of the production for X_i is explicitly performed. When $X_i \to uX_jvX_k$, then

$$\begin{aligned}\text{word}(X_i) &= u\,\text{word}(X_j)v\,\text{word}(X_k) \quad \text{and}\\ \text{word}(X'_i) &= PC_{ab\to c}(u)\,\text{word}(X'_j)PC_{ab\to c}(v)\,\text{word}(X'_k)\\ &= PC_{ab\to c}(u)PC_{ab\to c}(\text{word}(X'_j))PC_{ab\to c}(v)PC_{ab\to c}(\text{word}(X'_k)),\end{aligned}$$

with the last equality following by the induction assumption. Notice, that since ab is a non-crossing pair, all occurrences of ab in word(X_i) are contained in u, v, word(X_j) or word(X_k), as otherwise ab would be a crossing pair, which contradicts the assumption. Thus,

$$PC_{ab\to c}(\text{word}(X_i)) = PC_{ab\to c}(u)PC_{ab\to c}(\text{word}(X'_j))PC_{ab\to c}(v)PC_{ab\to c}(\text{word}(X'_k)),$$

which shows that $PC_{ab\to c}(\text{word}(X_i)) = \text{word}(X'_i)$, ending the proof of the claim. □

The second claim similarly establishes, how the pair compression of a non-crossing pair affects the NFA. To be more precise, what happens to a string defined by a path in the NFA after applying pair compression to the underlying NFA.

Claim 2. Consider a non-crossing pair ab and a path $\mathcal{P}$ in NFA N, which defines a list of labels:

$$u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n},$$

where each $u_{i_j} \in \Sigma^*$ is a string representing the consecutive letter labels and X_{i_j} represents a transition by a nonterminal transition. Then

$$PC_{ab\to c}(\mathrm{word}(\mathcal{P})) = PC_{ab\to c}(u_{i_1})\,\mathrm{word}(X'_{i_2})PC_{ab\to c}(u_{i_3})\,\mathrm{word}(X'_{i_4})\cdots\mathrm{word}(X'_{i_{n-1}})PC_{ab\to c}(u_{i_n}). \tag{3}$$

Proof. Similarly as in Claim 1, notice, that as ab is a non-crossing pair, the appearance of ab in the string defined by $\mathcal{P}$ cannot be split between a nonterminal and a string (other nonterminal). Thus, replacement of pairs ab takes place either wholly inside string u or inside $\mathrm{word}(X_i)$. The former is done explicitly by $PC_{ab\to c}$, while (2) establishes the form of the latter. This ends claim's proof. □

Now, after proving Claims 2–1, it is easy to show the main thesis of the lemma, i.e. that $\mathrm{word}(X'_n)$ is accepted by N' if and only if $\mathrm{word}(X_n)$ is accepted by N.

$\Leftarrow$ Suppose first that $\mathrm{word}(X_n)$ is accepted by N. Consider the accepting path $\mathcal{P}$ for $\mathrm{word}(X_n)$, represent it as a list of labels: $u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n}$, similarly as in Claim 2. Of course,

$$\mathrm{word}(\mathcal{P}) = \mathrm{word}(X_n) = u_{i_1}\,\mathrm{word}(X_{i_2})u_{i_3}\,\mathrm{word}(X_{i_4})\cdots\mathrm{word}(X_{i_{n-1}})u_{i_n}. \tag{4}$$

We will construct an accepting path $\mathcal{P}'$ in N' inducing a list of labels

$$PC_{ab\to c}(u_{i_1})X'_{i_2}PC_{ab\to c}(u_{i_3})X'_{i_4}\cdots X'_{i_{n-1}}PC_{ab\to c}(u_{i_n}). \tag{5}$$

Using (3) and recalling that $\mathrm{word}(\mathcal{P}) = \mathrm{word}(X_n)$ will be enough to conclude that $\mathrm{word}(\mathcal{P}') = PC_{ab\to c}(\mathrm{word}(X_n))$.

Notice, that by Algorithm 2

- if there is a transition $\delta_N(p, d, q)$ for a letter $d \in \Sigma$ in N, then there is the same transition $\delta_{N'}(p, d, q)$ in N'.
- if there is a path from p to q for a string ab in N then there is a transition $\delta_{N'}(p, c, q)$ in N'.

Thus, by a trivial induction on the length of the string u, if $\delta_N(p, u, q)$ then also $\delta_{N'}(p, PC_{ab\to c}(u), q)$. The situation is similar for nonterminals: if there is a transition $\delta_N(p, X_i, q)$ in N, then there is an analogous transition $\delta_{N'}(p, X'_i, q)$ in N'. Thus, a path $\mathcal{P}'$ with the same starting and ending state as $\mathcal{P}$ and the list of labels as in (5) is inductively defined. Since the starting (accepting) state in N and N' coincide, this shows that $\mathrm{word}(\mathcal{P}')$ is accepted by N'.

$\Rightarrow$ Suppose now that a string $PC_{ab\to c}(\mathrm{word}(X'_n))$ is recognised by N'. Let the path of the accepting computation in N' be $\mathcal{P}'$, with a list of labels

$$u'_{i_1} X'_{i_2} u'_{i_3} X'_{i_4} \cdots X'_{i_{n-1}} u'_{i_n}.$$

Similarly to the previous case, we will inductively define an accepting path $\mathcal{P}$ in N with a list of labels

$$u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n}, \tag{6}$$

where u_{i_j} is obtained from u'_{i_j} by replacing each c by ab.

Notice, that by Algorithm 2,

- if there is a letter transition $\delta_{N'}(p, c, q)$ in N', there is a path from p to q for a string ab in N.
- if the a letter transition $\delta_{N'}(p, d, q)$ for a letter $d \neq c$, there there is the same transition $\delta_N(p, d, q)$ in N.

Thus, by a simple induction, we conclude that if there is path in N' from p to q for a string u'_{i_j}, then there is a path in N from p to q for a string u_{i_j}. Now observe, that if there is a transition $\delta_{N'}(p, X_i, q)$ in N', then there is an analogous transition $\delta_N(p, X_i, q)$ in N. This completes the construction of $\mathcal{P}$. Since the starting and accepting states in N and N' coincide, the constructed path $\mathcal{P}$ is also accepting, furthermore, $\mathcal{P}$'s list of labels is as in (6).

It is left to show, that $\mathrm{word}(\mathcal{P}) = \mathrm{word}(X_n)$. Since $PC_{ab\to c}$ is a one-to-one function, it is enough to show that $PC_{ab\to c}(\mathrm{word}(\mathcal{P})) = PC_{ab\to c}(\mathrm{word}(X_n))$. Notice, that the latter equals $\mathrm{word}(X'_n)$, by (2).

Using (3) we can conclude that

$$\begin{aligned} PC_{ab\to c}(\mathrm{word}(\mathcal{P})) &= PC_{ab\to c}(u_{i_1})\,\mathrm{word}(X'_{i_2})PC_{ab\to c}(u_{i_3})\,\mathrm{word}(X'_{i_4})\cdots\mathrm{word}(X_{i_{n-1}}) \\ &= u'_{i_1}\,\mathrm{word}(X'_{i_2})u'_{i_3}\,\mathrm{word}(X'_{i_4})\cdots\mathrm{word}(X'_{i_{n-1}})u'_{i_n} \\ &= \mathrm{word}(\mathcal{P}') \\ &= \mathrm{word}(X'_n), \end{aligned}$$

which concludes the proof. □

(see page 8)

Proof of Lemma 6, stronger version. We shall proof Lemma 6 in a little stronger version: we allow the grammar to be in an a-succinct form and NFA N is assumed to satisfy the a-relaxed-(Aut 1), see appropriate definition between Algorithm 4 and Lemma 7. This weaker assumption will allow proving Lemma 8 here.

Notice, that by Lemma 12, there are only $|G| + 3n$ possible lengths of a's non-extendible appearance and that they can be calculated in polytime; hence line 1 takes polytime. All the loops in Algorithm 3 have only polynomially many iterations. All operations listed in Algorithm 3 are elementary and can be clearly performed in polytime, except replacing each a^ℓ by a_ℓ in a rule in line 4 and for the verification in line 6, for the former operation, it is enough to read the rules of the grammar: recall, that a^ℓ is represented as a pair (a, ℓ). Since $\ell \le 2^n$, addition of the lengths can be performed in polytime. The verification is more involved, we outline how to perform it: For given two states p, q and a string a^ℓ we want to verify, whether there is a path from p to q for a string a^ℓ (notice, that since a is inner, none of $\mathrm{word}(X_i)$ is a power of a). This can be rephrased as a fully compressed membership problem for NFA over a unary alphabet: it is enough to

- restrict NFA N to transitions by powers of a,
- make p the unique starting state,
- make q the unique accepting state.

Notice, that we can restrict ourselves by transitions by powers of a: since a is an inner letter, no $\mathrm{word}(X_i)$ begins or ends with a. Hence, when a path in the NFA defines a^ℓ, then this path uses only transitions by powers of a.

Observe, that all considered powers a^ℓ appear in strings defined by G, and so $\ell \le 2^n$., So each such a^ℓ can be represented by an SLP of polynomial size. In particular, Lemma 1 is applicable here and so the verification is in npolytime.

Concerning the preservation of invariants: we first show, that the G' is not in the a-succinct form, nor N' is a-relaxed. When Algorithm 3 finishes its work, all non-extendible appearances of a are replaced, in particular, there are no succinct representations of a powers inside the grammar. This should apply to transitions labelled with powers of a in N, and so the last line of Algorithm 3 should be modified, so that all transitions by *powers* of a are removed.

Now we can return to showing the preservation of the invariants: since the only change to the productions consists of replacing non-extendible occurrences of a by a single letter, (SLP 1)–(SLP 3) are preserved. Also, the only modifications to the NFA is the addition of new letter transitions. Thus, (Aut 1) holds. To see that also (Aut 2) holds, notice, that if p receives new incoming (outgoing) transition in N', this transition is of the form a_ℓ and p had an incoming

(outgoing, respectively) transition by a^ℓ in N. In particular, the starting and accepting state remain unaffected and no transition by \$ and # are introduced. Thus, also (Aut 2) holds for N'.

Notice, that if N is deterministic, so is N': suppose that there are two different transitions by d in N'. If d is not one of the new letters, i.e. it is a nonterminal or an old letter, then the same transitions were present in N, contradiction. If there are two transitions by a new letter a_ℓ from the same state p, then in N there are two different paths for a^ℓ from state p, contradiction.

Claim 3. After performing Algorithm 3, it holds that

$$\text{word}(X'_i) = AC_a(\text{word}(X_i)). \tag{7}$$

Claim 4. Consider an inner letter a and a path $\mathcal{P}$ in N, which has a list of labels:

$$u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n},$$

where each $u_{i_j} \in \Sigma^*$ is a string representing the consecutive letter labels and X_{i_j} represents a nonterminal transition, similarly as in Claim 2. Then

$$AC_a(\text{word}(\mathcal{P})) = AC_a(u_{i_1})\,\text{word}(X'_{i_2})AC_a(u_{i_3})\,\text{word}(X'_{i_4})\cdots\text{word}(X'_{i_{n-1}})AC_a(u_{i_n}). \tag{8}$$

The proofs are analogous as in Lemma 5 and are thus omitted. Notice, that the properties stated in Claims 3–4 do *not* depend on the non-deterministic choices of Algorithm 3.

It is left to show the main claims of the lemma: N recognises word(X_n) if and only if NFA N' recognises word(X'_n) for some non-deterministic choices.

$\Leftarrow$ Suppose first that the N' accepts word(X'_n), using the path $\mathcal{P}'$. Clearly word$(\mathcal{P}')$ = word(X'_n). Let the list of labels on $\mathcal{P}'$ be

$$u'_{i_1} X'_{i_2} u'_{i_3} X'_{i_4} \cdots X'_{i_{n-1}} u'_{i_n}.$$

Let u'_{i_j} be obtained from u_{i_j} be replacing each a_{ℓ_m} with a^{ℓ_m}. We shall construct a path $\mathcal{P}$ in N, which has the same starting and ending as $\mathcal{P}'$ and induces a list of labels

$$u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n}.$$

Notice, that

- if there is a transition $\delta_{N'}(p, X'_i, q)$ in N' then there is a transition $\delta_N(p, X_i, q)$ in N
- if there is a transition $\delta_{N'}(p, b, q)$ for $b \neq a$ in N', then there is a transition $\delta_N(p, b, q)$ in N
- if there is a transition $\delta_N(p, a_{\ell_m}, q)$ in N' for some a^{ℓ_m} that has a non-extendible in one of word(X_1), …, X_n, then there is a path from p to q for a string a^{ℓ_m} in N.

Therefore, by an easy induction, $\mathcal{P}$ is a valid path in N, moreover, since $\mathcal{P}'$ is accepting, so is $\mathcal{P}$. It is left to demonstrate, that $\mathcal{P}$ defines word(X_n): since AC_a is a one-to-one function, it is enough to show that $AC_a(\text{word}(\mathcal{P})) = AC_a(\text{word}(X_n))$. By (7) it holds that $AC_a(\text{word}(X_n)) = \text{word}(X'_n)$. The value of $AC_a(\text{word}(\mathcal{P}))$ is already known from (8), and so it is enough to show that

$$AC_a(u_{i_1})\,\text{word}(X'_{i_2})AC_a(u_{i_3})\,\text{word}(X'_{i_4})\cdots\text{word}(X'_{i_{n-1}})AC_a(u_{i_n}) = \text{word}(X'_n).$$

But this is simply the fact, that path $\mathcal{P}'$ defines word(X'_n), which hold by the assumption.

$\Rightarrow$ Suppose now that N accepts word(X_n). Consider the case in which Algorithm 3 always made a correct non-deterministic choice, i.e. that each time it correctly verified in line 6.

Let the accepting path $\mathcal{P}$ in N has a list of labels

$$u_{i_1} X_{i_2} u_{i_3} X_{i_4} \cdots X_{i_{n-1}} u_{i_n},$$

where, similarly as in Claim 4, each u_{i_j} is a string representing the consecutive letter labels (u_{i_j} may be empty) and X_{i_j} represents a transition by a nonterminal transition. By the definition,

$$\text{word}(X_n) = u_{i_1}\,\text{word}(X_{i_2})u_{i_3}\,\text{word}(X_{i_4})\cdots\text{word}(X_{i_{n-1}})u_{i_n}.$$

Now consider the a's appearance compression applied to both side of this equality. By (7) and (8)

$$\text{word}(X'_n) = AC_a(u_{i_1})\,\text{word}(X'_{i_2})AC_a(u_{i_3})\,\text{word}(X'_{i_4})\cdots\text{word}(X'_{i_{n-1}})AC_a(u_{i_n}).$$

We will construct an accepting path $\mathcal{P}'$ with a list of labels

$$AC_a(u_{i_1})X'_{i_2}AC_a(u_{i_3})X'_{i_4}\cdots X'_{i_{n-1}}AC_a(u_{i_n}).$$

Notice, that $\text{word}(\mathcal{P}') = \text{word}(X'_n)$, and so construction of such path $\mathcal{P}'$ will conclude the proof. We iteratively transform $\mathcal{P}$ into $\mathcal{P}'$. Notice, that

- if there is a transition $\delta_N(p, X_i, q)$ in N then there is a transition $\delta_{N'}(p, X'_i, q)$ in N'
- if there is a transition $\delta_N(p, b, q)$ for $b \neq a$ in N, then there is a transition $\delta_{N'}(p, b, q)$ in N'
- if there is a path in N from p to q for string a^ℓ that has non-extendible appearance in some $\text{word}(X_i)$, then there is a transition $\delta_{N'}(p, a_\ell, q)$ in N' (by the assumption that Algorithm 3 guessed correctly)

And by an easy induction $\mathcal{P}'$ is valid path in N' and has the same starting and ending state as $\mathcal{P}$. Since $\mathcal{P}$ is accepting in N and the starting and accepting states in N and N' coincide, $\mathcal{P}'$ is an accepting path in N'. □

(see page 9)

Proof of Lemma 7. We first explain, how to calculate the a-prefix (a-suffix); since G is in a-succinct form, this might be non-obvious. It is enough to scan the explicit strings stored in the productions' right-hand sides, summing the lengths of the consecutive a's appearances. This clearly works in polytime also for G stored in an a-succinct form, as the powers of a may have length at most 2^n, and so the length of their representation is linear in n. (the correctness of this approach is shown later).

Algorithm 4 runs in polytime, as the first loop has n iterations and the second $|N|$ iterations and each line can be performed in polytime.

Concerning the preservation of the invariants: in each rule of the grammar at most 4 non-extendible appearances of a are introduced. They may be long, however, in compressed form we treat them as singular symbols. In this way rules of G are stored in an a-succinct form, which was explicitly allowed. Then the a-prefix and suffix are removed and nonterminals defining ϵ are removed from the right-hand sides of the productions. This does not affect the (SLP 1)–(SLP 3) (recall, that a is not \$, neither #). Since the NFA is also changed, we inspect the invariants regarding N: introducing new states p_1, q_1 and replacing transition $\delta_N(p, X_i, q)$ by three transitions $\delta_{N'}(p, a^{\ell_i}, p_1)$, $\delta_{N'}(p_1, X, q_1)$, $\delta_{N'}(q_1, a^{r_i}, q)$ preserves the (Aut 1)–(Aut 2), with the exception that it a-relaxes (Aut 1).

Notice, that if N is deterministic, so is N': as already mentioned the only change done to N is the replacement of transition by X_i by a path of three transitions, such the the two states in the middle have exactly one incoming and outgoing transition, which clearly preserves determinism of the automaton.

To show the correctness, we prove by induction on i the two main claims of the lemma:

- Algorithm 4 correctly calculates the length of the a-prefix and a-suffix of $\text{word}(X_i)$, i.e. ℓ_i and r_i,
- $\text{word}(X_i) = a^{\ell_i}\,\text{word}(X'_i)a^{r_i}$.

As simple corollary, observe that these two conditions imply that a is not the last, nor the first letter of $\text{word}(X'_i)$.

For $i = 1$ notice, that the whole production for X_1 is stored explicitly, and so Algorithm 4 correctly calculates the a-prefix and suffix of $\text{word}(X_1)$ and after their removal, $\text{word}(X_1) = a^{\ell_1}\,\text{word}(X'_1)a^{r_1}$.

For the induction step, let $X_i \to uX_jvX_k$. Then by the induction assumption:

$$\begin{aligned}\operatorname{word}(X_i) &= u \operatorname{word}(X_j) v \operatorname{word}(X_k) \\ &= ua^{\ell_j} \operatorname{word}(X'_j) a^{r_j} v a^{\ell_k} \operatorname{word}(X'_k) a^{r_k}.\end{aligned}$$

There are cases to consider, depending on whether $\operatorname{word}(X'_j) = \epsilon$ or not and whether $\operatorname{word}(X'_k) = \epsilon$ or not. We will describe the one with $\operatorname{word}(X'_j) \neq \epsilon$ and $\operatorname{word}(X'_k) = \epsilon$, other cases are treated in a similar way. Then the rule is rewritten as $ua^{\ell_j} \operatorname{word}(X'_j) a^{r_j} v a^{\ell_k} a^{r_k}$. Since by the inductive assumption, a is not the first, nor the last letter of $\operatorname{word}(X'_j)$, the a prefix (a-suffix) of $\operatorname{word}(X'_j)$ is the a-prefix (a-suffix, respectively) of ua^{ℓ_j} (of $a^{r_j} v a^{\ell_k} a^{r_k}$, respectively). Thus it is correctly calculated by Algorithm 4. As the last action, Algorithm 4 removes the a-prefix and a-suffix, which shows that $a^{\ell_i} \operatorname{word}(X'_i) a^{r_i} = \operatorname{word}(X_i)$.

It is left to show that N accepts $\operatorname{word}(X_n)$ if and only if N' accepts $\operatorname{word}(X'_n)$. To this end, notice, that the only modification to N is the replacement of the transition of the form $\delta_N(p, X_i, q)$ by a path labelled with $a^{r_i}, X'_i, a^{\ell_i}$. Furthermore, the two vertices inside this path have only one incoming and one outgoing transition. The path labelled with $a^{r_i}, X'_i, a^{\ell_i}$ defines the string $a^{r_i} \operatorname{word}(X'_i) a^{\ell_i}$, which was already shown to be $\operatorname{word}(X_i)$. It is left to observe, that none of the newly introduced states in the middle of the path is accepting, nor starting. Hence the starting (accepting) states of N and N' coincide, each string is accepted by N if and only if it is accepted by N'. □

(see page 9)

Proof of Lemma 8. Notice, that the proof of Lemma 6 was shown in a stronger version, which coincides with the statement of Lemma 8. □

(see page 10)

Proof of Lemma 9. Majority of the proof is similar to the proof of Lemma 7. Nevertheless, it is written for completeness.

The loop is executed polynomially many times, and also each line of the code can be performed in polytime, and so in total Algorithm 5 runs in polytime.

Concerning the preservation of invariants: since the only operation performed on G is replacing nonterminal X_i by aX'_i and then deleting the first letter of the nonterminal, and nonterminals generating ϵ are explicitly removed from the rules, the resulting grammar is in the form (1). Notice, that the invariants (SLP 1)–(SLP 2) are clearly preserved by the listed operations. Since the first and last symbol of $\operatorname{word}(X_n)$ is not modified, (SLP 3) holds as well.

Let us move to the NFA invariants: the only change applied to NFA is the replacement of the transitions $\delta_N(p, X_i, q)$ to by a path $\delta_{N'}(p, \operatorname{first}(X_i), p_1)$, $\delta_{N'}(p_1, X_i, q)$, or by $\delta_{N'}(p, \operatorname{first}(X_i), q)$, which does not affect (Aut 1)–(Aut 2). For the same reason, if N is deterministic, so is N':

We first show by induction the second claim of the lemma, i.e. that if $\operatorname{word}(X_i) = bu$ for some $b \in \Sigma$, $u \in \Sigma^*$, then $\operatorname{word}(X'_i) = u$. For the induction basis consider $i = 1$: let the rule for X_1 be $X_1 \to bu$ for some $b \in \Sigma$ and $u \in \Sigma^*$. Then $x_1 = b$ and $u' = u$ and the production for X'_1 is $X'_1 \to u'$. Clearly the claim holds in this case.

For the inductive step, consider $i > 1$ and let X_i's production be $X_i \to uX_jvX_k$. Then, $\operatorname{word}(X_i) = u \operatorname{word}(X_j) v \operatorname{word}(X_k)$ and by by induction assumption:

$$u \operatorname{word}(X_j) v \operatorname{word}(X_k) = u \operatorname{first}(X_j) \operatorname{word}(X'_j) v \operatorname{first}(X_k) \operatorname{word}(X'_k).$$

Since $u \operatorname{first}(X_j)$ is non-empty, $\operatorname{first}(X_i)$ can be computed in line 4 of Algorithm 5. We conclude that

$$\begin{aligned}\operatorname{word}(X_i) &= u \operatorname{first}(X_j) \operatorname{word}(X'_j) v \operatorname{first}(X_k) \operatorname{word}(X'_k) \\ &= \operatorname{first}(X_i) u' \operatorname{first}(X_j) \operatorname{word}(X'_j) v \operatorname{first}(X_k) \operatorname{word}(X'_k) \\ &= \operatorname{first}(X_i) \operatorname{word}(X'_i).\end{aligned}$$

The analysis for other forms of the rule, i.e. $X_i \to uX_jv$ or $X_i \to u$, is similar.

We show now that N' accepts exactly the same strings as N. The only change done in the NFA is the replacement of transitions of the form $\delta_N(p, X_i, q)$ by a path inducing list of labels with

first(X_i) and X'_i or by a transition $\delta_N(p, \text{first}(X_i), q)$, let us consider the former case, the latter is similar. Notice, that word(X_i) = first(X_i) word(X'_i) and so the new path denotes the same string, as the replaced transition. Furthermore, the newly introduced state in the middle of this path has only one ingoing and outgoing transition. The starting and accepting states were not modified, and so the both automata recognise the same strings. Since word(X_n) = word(X'_n), this ends the proof.

It is left to show, that none of the pairs $a_\ell b$ is crossing. Assume for the sake of contradiction, that some of such pairs $a_\ell b$ is. The analysis splits between cases, why the pair $a_\ell b$ is crossing.

$a_\ell b$ has a crossing appearance in the NFA N' In this case there is a state p and a pair of incoming and outgoing transitions, such that a_ℓ is the last letter of the string encoded on the incoming transition and b is the first letter on the outgoing string encoded transition. Moreover, at least one of this transitions is labelled with a nonterminal, we distinguish these two subcases:

- **transition into p is labelled with X'_i** In such a case a_ℓ is the last letter of word(X'_i), and since word(X_i) = first(X_i) word(X'_i), also the last letter of word(X_i). Thus a_ℓ is a right-outer letter in G, contradiction with Lemma 8.
- **transition into p is labelled with a letter and transition from p is labelled with X'_i** In this case a_ℓ is the label of the transition incoming to p. Observe, that due to the algorithm, if X'_i labels the transmission from p, there is a unique transition to p, labelled with first(X_i). This means that a_ℓ is the first letter of word(X_i), which is not possible by Lemma 8.

$a_\ell b$ has a crossing appearance in nonterminal X'_i Let the production for X'_i be $X'_i \to uX'_jvX'_k$, the cases of other productions are similar. There are three possibilities for a crossing appearance of $a_\ell b$ in word(X'_i):

- **a_ℓ is the last letter in u and b is the first letter in** word(X'_j) Algorithm 5 replaced X_j with first(X_j)X'_j in the rule for X'_i. Thus, since a_ℓ is the letter preceding X'_j in the rule for X'_i, it holds that a_ℓ was the first letter in word(X_j), which is a contradiction with Lemma 8.
- **a_ℓ is the last letter in** word(X'_i) **and b is the first letter in v** Since the last letter of word(X'_i) and word(X_i) are the same, as word(X_i) = first(X_i) word(X'_i), a_ℓ is the last letter of word(X_i), which is a contradiction with Lemma 8.
- **a_ℓ is the last letter in v and b is the first letter in** word(X'_k) the proof is the same as in the first case. □

Proof of Lemma 10. Notice first, that there are at most $|G|$ different letters a_i: they were introduced by Algorithm 3 as representations of non-extendible appearances of a, and by Lemma 12 there are at most $|G|$ different such appearances. Other operations of Algorithm 6 constitute of call to Algorithm 5 and Algorithm 2, and thus the running time and preservation of the invariants follows properties of these two algorithms, stated in Lemma 8 and Lemma 9 □ (see page 10)

Proof of Lemma 11. We first bound the size of Σ. We show, that at the beginning of each iteration of the main loop in Algorithm 1 each right-hand side of the production has at most $48n$ explicit letters, and that inside each iteration of the main loop of Algorith 1 there are at most $12n$ new letters added (we exclude the letter replacing compressed strings). Notice, that during the run of Algorithm 1 grammar G may be in succinct form, and accordingly we treat a^ℓ as one symbol. (see page 10)

Clearly, these bounds are true when Algorithm 1 starts working. Let us fix a rule and consider, how many new letters may be introduced in this rule. There are only two cases, in which new letters (not coming from compression) are introduced to a rule:

changing an outer letter to an inner one (line 3 of Algorithm 4) There are at most 4 new powers of a (all possibly in succinct form) that may be introduced in each invocation of Algorithm 4. While these powers of a are written in a succinct representation, they will be all replaced by single letters in the later appearance compression for a.

In total, Algorithm 4 is invoked for each outer letter, and there are at most $2n$ such letters, by Lemma 2. So there are at most $8n$ new symbols introduced in this way for one iteration of the main loop of Algorithm 1.

popping letters (line 3 of Algorithm 5) Each invocation of of Algorithm 5 introduces at most 2 new symbols. As Algorithm 5 is also used for each outer letter, in this way at most $4n$ new symbols were introduced per each iteration of the main loop of Algorithm 1.

Hence, there are at most $12n$ new letters added to the right-hand side of a production in each iteration of the main loop in Algorithm 1. Still, the main task performed by Algorithm 1 is the compression: an argument similar as in the proof Lemma 4 can be used to show that the size of the explicit strings in the rules decreases by a factor of 3/4 in each iteration of loop from line 1 in Algorithm 1. Of course, the newly introduced letters may be unaffected by this compression. It is left to verify, that $48n$ is indeed the upper bound on the size of the right-hand size of a rule, let it be $X_i \to \alpha_i$:

$$|\alpha_i'| \leq \frac{3}{4} \cdot |\alpha_i| + 12n \leq \frac{3}{4} \cdot 48n + 12n \leq 48n.$$

Which proves the bound on $|G|$ at the end of each iteration of the main loop. Notice, that as there at most $12n$ new letters added to this rule, inside the phase the size of G is at most $60n^2$.

We now turn our attention to the size of Σ. Again, consider the execution of Algorithm 1 and one iteration of the main loop. We show, that there are $\mathsf{poly}(n)$ letters introduced in one such iteration. New letters are introduced when compression of pairs or appearance compression is applied. There are the following possibilities

compression of an inner letter (line 1 of Algorithm 2) Each compression of an inner letter decreases the total length of explicit strings used in G by at least 1. Since the size of each right-hand side is at most $48n$ at the beginning of the iteration and there are at most $12n$ new letters introduced to a rule in each iteration, there can be at most $60n^2$ such compressions.

compression of a non-crossing pair (line 4 of Algorithm 3) The same argument as in the previous case applies.

appearance compression for of an outer letter (Algorithm 4 followed by Algorithm 3) There are at most $2n$ outer letters, by Lemma 2, and each of them has at most $|G|$ different lengths of non-extendible appearances, by Lemma 12. Compressing all of them introduces at most $2n|G|$ new letters to Σ.

compression of a crossing pair (call to Algorithm 2 made by Algorithm 6) The crossing pairs compression is run for each of the $2n$ outer letters. When such an outer letter is fixed, there are at most $|G|+3n$ different pairs of letters appearing in $\mathrm{word}(X_n)$, by Lemma 2, this may introduce up to $|G|+3n$ new letters. Hence, in total this might introduce up $2n|G|+6n^2$ new letters to Σ.

Thus, the size of $|\Sigma|$ is $\mathsf{poly}(n)$.

It is left to bound the size of $|N|$: as there are at most n non-terminal transitions, the the size of transition function of N is at most $\mathcal{O}(|Q_N|^2 \cdot (|\Sigma| + n))$. So it enough to bound the number of states of N by $\mathsf{poly}(n)$. Recall, that for simplicity we assumed that all transitions in the input DFA are labelled with different SLPs, in particular, this polynomial bound holds for the input instance.

There are only two situations, in which new states are introduced:

changing a from outer letter to inner one (Algorithm 4 in line 6) Algorithm 4 introduces two states per nonterminal transition, and there are at most n such transitions, by (Aut 1). So it is enough to estimate, how many times Algorithm 4 is invoked. Algorithm 4 is run for each outer letter, and there are at most $2n$ outer letters. Thus, in one iteration of the main loop of Algorithm 1 it adds at most $\mathcal{O}(n^2)$ states in total in this case.

popping letters (Algorithm 5 in line 10) A similar argument applies, with the only difference, that Algorithm 5 introduces one new state per nonterminal transition.

It is left to recall, that by Lemma 4 the main loop of Algorithm 1 is run $\mathcal{O}(n)$ times. This ends the proof of the lemma. □